(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number : 0 455 402 A2

(12) # EUROPEAN PATENT APPLICATION

(21) Application number : 91303643.0

(22) Date of filing : 23.04.91

(51) Int. Cl.[5] : H04L 12/24

(30) Priority : 03.05.90 US 519187

(43) Date of publication of application :
06.11.91 Bulletin 91/45

(84) Designated Contracting States :
DE FR GB

(71) Applicant : Hewlett-Packard Company
Mail Stop 20 B-O, 3000 Hanover Street
Palo Alto, California 94304 (US)

(72) Inventor : Wu, Jeff C.
2630 Wapiti Road
Fort Collins, Colorado 80525 (US)

(74) Representative : Colgan, Stephen James et al
CARPMAELS & RANSFORD 43 Bloomsbury Square
London WC1A 2RA (GB)

(54) **Automatic discovery of network elements.**

(57) Disclosed is a computer network node discovery system that provides a general way of discovering network elements, or nodes, connected to a computer network, and a specific algorithm for discovering nodes connected to a TCP/IP network, using the SNMP protocol available within the TCP/IP network software. Some nodes on a network, called discovery agents, can convey knowledge of the existence of other nodes on the network. The network discovery system queries these agents and obtains the information they have about other nodes on the network. It then queries each of the nodes obtained to determine if that node is also a discovery agent. In this manner, most of the nodes on a network can be discovered. The process of querying discovery agents to obtain a list of nodes known to the discovery agents is repeated at timed intervals to obtain information about nodes that are not always active. In a TCP/IP network, discovery agents are nodes that respond to queries for an address translation table which translates internet protocol (IP) addresses to physical addresses. The data from each node's address translation table is used to obtain both the IP and the physical address of other nodes on the network. These nodes are then queried to obtain additional information. After all the nodes on a network are discovered, the list of nodes is written to a database where it can be displayed by the network manager or other users of the network.

EP 0 455 402 A2

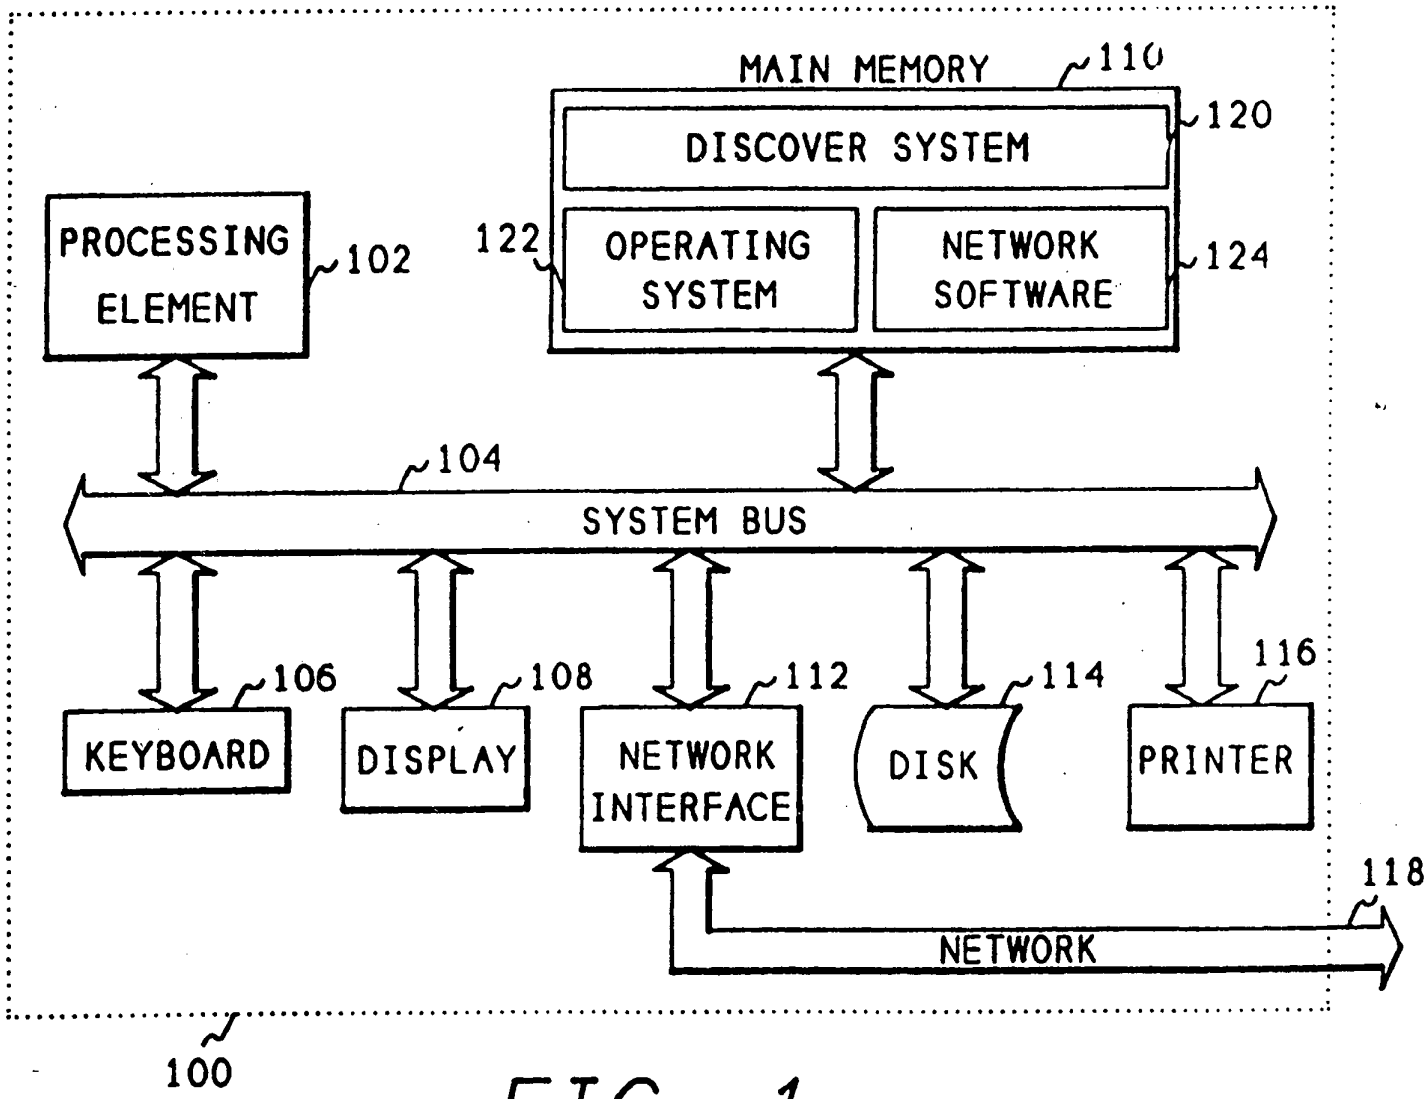
MAIN MEMORY
110
DISCOVER SYSTEM
120
PROCESSING ELEMENT
102
122
OPERATING SYSTEM
NETWORK SOFTWARE
124
104
SYSTEM BUS
106
108
112
114
116
KEYBOARD
DISPLAY
NETWORK INTERFACE
DISK
PRINTER
118
NETWORK
100

FIG. 1

FIELD OF THE INVENTION

This invention relates to computer systems and more particularly to computer networks that interconnect computers. Even more particularly, the invention relates to determining the nodes connected to a network.

BACKGROUND OF THE INVENTION

Computer networks are collections of hardware and software that connect computers and allow them to send information from one computer to another electronically. A computer network is comprised of the physical hardware connections between the various computers, for example telephone lines or a coax cable, and the software used to send and receive data and to route the data to the selected computer on the network.

A local area network (LAN) is a network connection between computers in close proximity, typically less than one mile, and usually connected by a single cable such as coax cable. A wide area network (WAN) is a network of computers located at longer distances, often connected by telephone lines or satellite links. Network software may sometimes be used with both types of networks. For example, a popular network is the Department of Defense internetworking protocol suite, known as Transmission Control Protocol/Internet Protocol (TCP/IP). This system was originally developed by the Defense Advanced Research Projects Agency (DARPA) and has now been widely distributed to Universities and industry.

When a network is fast growing, that is, network elements or nodes are being added frequently, a network administrator may not know all of the nodes connected to the network. Also, a network administrator new to his or her job may not be familiar with the nodes on the network. Determining the nodes manually is a difficult problem. The administrator may contact all the users of the network known to the administrator, however, infrequent users may be forgotten and not contacted. Also, if a node is connected to the network, but not active because the computer is not powered up or is inoperative, that node may not be included in the list. In a very short local area network, a network administrator may physically trace the cable of the network to determine which nodes are located on the network. However, since longer local area networks can extend as far as a mile, through many floors and offices within a building, physical tracing may be impossible. In a wide area network, physical tracing is almost always impossible.

For some commonly used networks, special equipment can be purchased that will determine the nodes located on the network and the distance between them. This equipment, called a probe, is often limited by the other components of the network, however. For example, in a local area network, a repeater unit may be used to extend the effective distance of the local area network to a distance greater than is capable with a single cable. A repeater unit amplifies signals, and therefore will not allow a probe to determine the location of nodes beyond the repeater.

Other units connected to the network may obscure nodes. For example a bridge unit connects two similar networks but only passes messages that are being sent from a node on one side of the bridge to a node on the other side of the bridge. It will not pass messages between nodes on the same side, in order to reduce the traffic on the other side of the bridge. A bridge will prevent a probe from determining the nodes on the other side of the bridge. A gateway is a unit that connects dissimilar networks to pass messages. Because a gateway may have to reformat a message to accommodate a different network protocol, it will prevent a probe from finding nodes beyond the gateway.

There is need in the art then for a method of determining the nodes on a local area network. There is further need in the art for determining such nodes without the use of special equipment. A still further need is for a method that will determine which nodes are located beyond the repeater units, bridges, and gateways on a network.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide a method of determining the elements or nodes connected to a network.

It is another object of the invention to provide a method of discovering network nodes on a TCP/IP network.

Another object of the invention is to determine which discovered nodes are discovery agents and can convey knowledge of the existence of other nodes on the network.

Another object is to query all discovery agents and ask for other nodes on the network

A further object is to query all TCP/IP nodes to retrieve the address translation table from the TCP/IP node.

The above and other objects of the invention are accomplished in a system which provides a general way of discovering network elements, or nodes, and a specific algorithm for discovering nodes within a TCP/IP network, using a standard Simple Network Management Protocol (SNMP), which is available within the TCP/IP network.

Some nodes on a network can convey knowledge of the existence of other nodes on a network, and are called discovery agents. When a network contains discovery agents, these agents can be queried to obtain the information they have about other nodes on the network. By obtaining a list of nodes from a single

discovery agent, and querying each of the nodes obtained to determine if it is also a discovery agent, most of the nodes on a network can be discovered.

The process of querying discovery agents to obtain a list of nodes known to be discovery agents, must be repeated at timed intervals. At any given time on a network, one or more nodes may not be responding to the network, either because it is inoperative, or because it is not powered up. Therefore, if the discovery process is attempted during this time, these unavailable nodes will not be discovered. By repeating the discovery process over time at regular intervals, additional nodes on a network can be discovered.

In a TCP/IP network, discovery agents are nodes that respond to queries for an address translation table. Within TCP/IP network, every node will have an internet protocol (IP) address. This address is a 32 bit number and is unique to all nodes within the TCP/IP network. Although the IP address is probably unique to all nodes everywhere that use the TCP/IP protocol, the physical address of a node on a particular network will be different from the IP address. For example, some types of LANs use an 8 bit address, and can therefore use the low order 8 bits of the IP address, however, some other types of LANs use a 48 bit address and cannot use the internet address. Therefore, every node within a TCP/IP network must have an address translation table which translates the IP address to the physical address. The data from each node's address translation table can be used to obtain both the IP and the physical address of other nodes on the network. Again, as described in the above general algorithm, the queries should be repeated at timed intervals to insure that recently activated nodes are discovered. Another reason for repeating the discovery process over timed intervals in a TCP/IP network is that some of the information within a node's address translation table may be purged if the node does not use the information after a period of time. This purge is used to reduce the table size requirements within a node. By repeating the queries at timed intervals, the greatest amount of translation table information may be obtained.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other objects, features, and advantages of the invention will be better understood by reading the following more particular description of the invention, presented in conjunction with the following drawings, wherein:

Fig. 1 shows a block diagram of the hardware of the node that runs the process of the present invention;

Fig. 2 shows a diagram of a typical computer interconnection network;

modules of the discovery system of the present invention;

Fig. 6 shows a flowchart of the main module of the invention;

Fig. 7 shows a flowchart of the self-seed module of the invention;

Fig. 8 shows a flowchart of the process-node module of the invention;

Fig. 9 shows a flowchart of the process-ping module of the invention;

Fig. 10 shows a flowchart of the process-IFIP module of the invention;

Fig. 11 shows a flowchart of the store-IP module of the invention;

Fig. 12 shows a flowchart of the store-IF module of the invention;

Fig. 13 shows a flowchart of the invalidnode module of the invention;

Fig. 14 shows a flowchart of the findnode module of the invention;

Fig. 15 shows a flowchart of the addnode module of the invention;

Fig. 16 shows a flowchart of the process-AT module of the invention; and

Fig. 17 shows a flowchart of the store-AT module of the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The following description is of the best presently contemplated mode of carrying out the present invention. This description is not to be taken in a limiting sense but is made merely for the purpose of describing the general principles of the invention. The scope of the invention should be determined by referencing the appended claims.

Fig. 1 shows a block diagram of the computer hardware that contains the discovery system of the present invention. Referring now to Fig. 1, a computer system 100 contains a processing element 102. The processing element 102 communicates to other elements within the computer system 100 over a system bus 104. A keyboard 106 is used to input information from a user of the system, and a display 108 is used to output information to the user. A network interface 112 is used to interface the system 100 to a network 118 to allow the computer system 100 to act as a node on a network. A disk 114 is used to store the software of the discovery system of the present invention, as well as to store the data base collected by the discovery system. A printer 116 can be used to provide a hard copy output of the nodes of the network discovered by the discovery system. A main memory 110 within the system 100 contains the discovery system 120 of the present invention. The discovery system 120 communicates with in operating system 122 and

network 118.

Fig. 2 shows a diagram of a network. Referring now to Fig. 2, a network 202 contains a node 206. Node 206 contains the processor 100 (Fig. 1) which contains the discovery system software of the present invention. Node 206 is attached to a first network segment 118. The network segment 118 is connected to a repeater 212 which is connected to a second network sequent 214. This second network system 214 has nodes 216 and 218 attached to it. A repeater, such as repeater 212, allows network sequents to be connected to allow a network to be extended over a longer distance. An important characteristic of a repeater is that there is no translation of data passing through it. That is, every message that is transmitted on one network segment, will pass unchanged through a repeater to the other network segment. Therefore, any messages broadcast, for example, by node 206 will be received by node 216 and node 218 after these messages pass through repeater 212.

Network segment 118 is also attached to a bridge 208 which connects it to a third network sequent 210. A bridge will only pass messages that are being transmitted from a node on one side of the bridge to a node on the other side of the bridge. It will block messages that are transmitted from a node on one side of the bridge to a node on that same side of the bridge. This characteristic reduces network traffic on various sequents of a network.

Segment 118 is also attached to a router/gateway 220 which connects is to a fourth network segment 222. Routers are devices that connect network segments which have similar characteristics. Gateways are devices which connect networks having different types of characteristics. For example, a gateway might connect a local area network to a wide area network.

Because bridges, routers, and gateways, must process the messages sent over the network, they also must contain information about which nodes are on the network. Therefore, bridges, routers, and gateways are authoritative sources of information for determining the nodes on the network. A protocol defines the format of messages that are sent across a network. One popular protocol is the Department of Defense Internetworking Protocol Suite, popularly known as TCP/IP. Because it was developed by the Department of Defense, this protocol is widely available and used extensively, particularly in a university environment. Also, this suite of protocols is very popular on the UNIX operating system and has seen wide distribution there. The internet protocol (IP) uses a single thirty-two bit address for all nodes that can be connected to the internet at any location. Physical addresses within a particular type of network, are normally different from an IP address. If a network address is very small, perhaps eight bits, it may be the same as the low order eight bits of the IP address. If a network address is large, for example, some LANs use forty-eight bit addresses, it is impossible for these addresses to correspond directly to IP addresses. Therefore, both an IP address and a physical address exist for each node on a network. Devices such as routers, gateways, and bridges, which can send messages from one network to another must be able to translate between IP addresses and physical addresses. Therefore, these devices have translation tables which allow them to translate between these two types of addresses. By accessing these translation tables, one of the nodes on a network can obtain information about the other nodes on the network. The existence of these translation tables allow the method of the present invention to perform its function.

A network probe 224 is also attached to the network 118. A network probe 224 is a device that assists in locating defective nodes and assists in repairing those nodes. Since it is a testing device, it may or may not be attached to a network at any given time. When a probe is attached to a network, the discovery system of the present invention can query the probe and use information obtained from the probe to assist in discovering other nodes on the network.

Figs. 3 through 5 show a hierarchy diagram of the modules of the software of the present invention. Referring now to Figs. 3 through 5, discovery module 302 is the main module of the system. Discovery calls self-seed block 304 to start the process of building a database about the network, and it calls process-node block 306 to process information about each node that it obtained from self-seed. Process-node block 306 calls process-ping block 308 to query a node on the network to determine if that node is active. Process-node block 306 also calls process-IFIP block 310 for each IP address that it obtains. Process-IFIP block 310 calls store-IP block 402 for each IP address, and store-IP block 402 calls invalidnode block 406, findnode block 408, and addnode block 410, for each IP address. For each IF entry (physical address) received, process-IFIP block 310 calls store-IF block 404. For each address translation table entry, process-node block 306 calls process-AT block 312 which in turn calls store-AT block 502. Store-AT block 502 calls invalidnode at block 504, findnode block 506, and addnode block 508.

Fig. 6 shows a flowchart of the discovery module block 302 (Fig. 3). Referring now to Fig. 6, after entry block 602 gets any options that the user wishes to enter. Block 604 then initializes the database used to permanently store the nodes, and loads node list from existing entries in the database. If a database for the network does not exist, the discovery system has the ability to create that database. If a database of the network already exists, the discovery system will use the node information which is already available in that database to query other nodes within the system.

Block 606 then initializes domains. A domain

defines the limit beyond which the user of the discovery system does not wish to find nodes. That is, the domain limits the range of the discovery process. This limitation is necessary on large networks, to keep the amount of processing to reasonable level. Furthermore, a user usually is only interested in the nodes on a particular network segment, or the network segment connected by repeaters and possibly bridges.

Block 608 then calls Fig. 7 to self-seed the system. If no entries were available in the database, the discovery system can self-seed by sending a broadcast message and determine who responds to that message. After returning from self-seed, block 610 points to the first node list entry. As discussed earlier, the node list will contain a list of the nodes already known to the system. This list can be input from the database, or the list can be started from self-seed module. After pointing to the first entry, block 612 determines if there are more entries to process. If there are no more entries to process, block 612 transfers to block 614 which will wait a predetermined period of time before reprocessing the entire node list. Typically, block 614 will wait for approximately thirty seconds. By reprocessing the node list periodically, additional nodes can be discovered. This is because a node may be inactive on the system at any given time and might not be discovered by a single pass through the network. By waiting and reprocessing the node list, nodes that were inactive may now be active and additional information can be obtained.

If more entries in the node list exist, block 612 transfers to block 616 to process one of the nodes. After processing that node, block 616 transfers to block 618 which points to the next node list entry and returns to block 612 to process the next node.

Fig. 7 shows a flowchart for the self-seed block 304 (Fig. 3) which obtains initial information about nodes on the network. Referring now to Fig. 7, after entry, block 702 sends an SP broadcast request to all nodes on the network. SNMP stands for Simple Network Management Protocol, and is a part of the TCP/IP network software. After sending the broadcast request, block 702 transfers to block 704 which receives SNMP messages from the nodes. If more SNMP messages are available, block 704 transfers to block 706 which adds a node to the node list for each message received. In this manner, all nodes that are currently active on the network can be queried to obtain initial information about the node. After all SNMP messages have been received, block 704 returns to the caller.

Another way of self-seeding is to query the address translation table for the node that is executing the discovery system. This table will contain the addresses of other nodes on the network, and these addresses are then used to start the discovery process.

(Fig. 3). The process-node module of Fig. 8 is called from the discovery module of Fig. 6 once for each entry in the node list. Therefore, whin Fig. 8 is called, the address of a single node is passed to it. Referring now to Fig. 8, after entry, block 802 determines whether the node is within a domain. As discussed earlier, the domain defines the limits beyond which the discovery program does not wish to discover new nodes. If the node is within the domain, block 802 transfers to block 804 which calls the process-ping module of Fig. 9 to determine whether the node is active. After returning from Fig. 9, block 804 transfers to block 806 to determine whether the state of the node has changed since the last information was obtained. That is, when the process-ping module queries the node, it determines the state of the node at the present time. This state is compared, in block 806, with the state of the node as it was known previously in the database. If that state has changed, block 806 transfers to block 808 to store the new state in the database. Control then returns to block 810 which calls process-IFIP to retrieve the IF and IP tables from the node. After returning from Fig. 10, block 810 transfers to block 812 which determines whether the node responded to an SNMP request. If the node did respond to the SNMP request, block 812 transfers to block 814 which determines whether the node is currently in the database. If the node is not in the database, block 814 transfers to block 816 to add the node to the database. Control then continues at block 818 which calls Fig. 16 to retrieve the address translation table from the node. Control then returns to the caller.

Fig. 9 shows a flowchart of the process-ping module block 308 (Fig. 3). This module is called to determine whether a node is active on the network. Referring now to Fig. 9, after entry block 902 determines whether the ping interval has elapsed. The ping interval is used to prevent a node from being queried too often. If the ping interval has not elapsed, block 902 returns to the caller. If the ping interval has elapsed, block 902 transfers to block 904 which sends an ICMP-echo message to the node. The ICMP-echo protocol is defined as a part of TCP/IP and is used to cause the node to return an acknowledgement to a message. Block 904 then transfers to block 906 which determines whether a response has been received from the other node. If a response has not been received within a predetermined amount of time, typically block 906 transfers to block 910 which sets a flag to indicates that the node failed to respond. If the node does respond, block 906 transfers to block 908 which sets a flag to indicate that the node did respond and then block 912 sets a new ping interval which will prevent the node from being pinged for the period of the interval. The ping interval is typically five minutes. Block 912 then returns to the caller.

Fig. 10 shows a flowchart of the process-IFIP

available in a node to define the translation of physical addresses to IP addresses. The information is available as two different tables, with an index contained in the IF table to cross-reference to the IP table within the node. By obtaining these two tables, the discovery system can determine what the other interfaces to which a node is connected, and therefore determine other networks to which the node is connected. Referring now to Fig. 10, after entry, block 1002 determines whether the IFIP interval has elapsed. The IFIP interval is similar to the ping interval described with respect to Fig. 9, and is used to keep a node from being queried too often. If the IFIP interval has not elapsed, block 1002 returns to the caller. If the IFIP has elapsed, block 1002 transfers to block 1004 which sends an SNMP message to request the node to send its next IP table entry to the discovery node. When an entry is received, block 1006 calls store-IP module of Fig. 11 to store the node within the node list. Block 1007 then transfers back to block 1004 if more IP entries are available. After all the entries are all stored in the node list, block 1007 transfers to block 1008 which sets a new IFIP interval of typically greater than 10 hours. Block 1010 then sends an SNMP message to request that the node send its next IF table entry to the discovery node. When an IF table entry is received, block 1012 calls the store-IF module of Fig. 12. Block 1014 then transfers back to block 101 if more entries are available. After receiving and storing all the IF table entries, block 1014 returns to the caller. Each IF table entry contains an index into the IP table. By using this index, physical addresses in the IF table can be matched with the IP address.

Fig. 11 shows a flowchart of the store-IP process block 402 (Fig. 4). Referring now to Fig. 11, after entry block 1102 calls Fig. 14 to find the node in the node list. The node will be found if the discovery system has already encountered this node in its process. Block 1304 then determines whether the node exists, and if the node does not exist, block 1104 transfers to block 1106 which calls Fig. 13 to determine whether the node is valid. Block 1108 then determines if the node is valid and if it is valid, block 1108 transfers to block 1110 to add the node to the node list. After adding the node, or if the node already existed, control goes to block 1112 which updates the state information about the node. After updating the node state information or if the node was not valid, Fig. 11 returns to the caller.

Fig. 12 is a flowchart of the store-IF process of block 404 (Fig. 4). This module is called for each table entry in the IF table received from a node. Referring now Fig. 12, after entry, block 1202 finds the IP index within the IF record. As described earlier, each IF table entry will have a corresponding IP table entry, and the IP entry is referenced by an index value contained in the IF entry. Block 1204 then determines whether a matching IP record exists. If a matching IP record does exist, block 1204 transfers to block 1206 which moves the physical address from the IP record to the node record in the node list. Block 1208 then updates any state information in the node record. After updating the state information, or if there were no matching IP record, Fig. 12 returns to its caller.

Fig. 13 shows a flowchart of the invalidnode module block 406 (Fig. 4). Referring now Fig. 13, after entry, block 1302 determines whether the address of the node is simply the loopback address of another node. Each node has a loopback address associated with it for use in testing the node. Because the loopback address refers to the same node, no additional information can be obtained from that node and the loopback address is never stored as a node address. If the IP address is not equal to the loopback address, block 1302 transfers to block 1304 to determine whether the node is within the domain. As described earlier, the domain is used to determine the limits beyond which the discovery system will not attempt to discover new nodes. If the node is within the domain, block 1304 transfers to block 1306 which returns an indication that the node is valid. If the node is not within the domain or if the IP address equals the loopback address, control transfers to block 1308 which returns an error indication indicating that node is not valid. Control then returns to the caller.

Fig. 14 is a flowchart of the findnode module block 408 (Fig. 4). The moduie is used to find a node within the node list. Referring now Fig. 14, after entry, block 1402 gets the node list entry. Block 1404 then determines whether the IP address matches the entry in the list. If a match does occur, block 1404 transfers to block 1408 which returns an indication that the node is in the node list. If the IP address does not match, block 1404 transfers to block 1406 which gets the next node list entry and block 1410 then determines whether the end of table has been reached. If the end of the list has not been reached, block 1410 transfers back to block 1404 to check the entry just found. If the end of the list has occurred, block 1410 transfers to block 1412 which returns an error indication indicating that the node is not in the node list.

Fig. 15 shows a flowchart of the process of adding a node to the node list. Referring now to Fig. 15, after entry, block 1502 performs a hash operation on the IP address to create a pointer into the node list. Block 1504 then allocates memory for a node record, and block 1506 stores the data available for the node into the node record at the location pointed to by the hashed IP address. Block 1506 then returns to the caller.

Fig. 16 shows a flowchart of the process-AT module of block 312 (Fig. 3). This module is called by the process-node module for each entry in the node list. Referring now to Fig. 16, after entry, block 1602 determines whether the AT interval has expired. The AT interval is used to prevent a node from being polled too frequently. If the AT interval has not expired, block

1602 simply returns to the caller. If the AT interval has expired, block 1602 transfers to block 1604 which sends an SNMP message to request that the node send its next address translation table entry to the discovery node. When an entry is received, block 1606 is called to store the table entry. Block 1607 then transfers back to block 1604 if more table entries are available. After storing all the table entries, block 1607 transfers to block 1608 which updates the node's state information in the node list. Block 1610 then sets a new AT interval, typically fifteen seconds, and returns to the caller.

Fig. 17 shows a flowchart of the store-AT module of block 502 (Fig. 5). Referring now to Fig. 17, after entry, block 1702 calls the findnode module Fig. 14 to determine whether the node is already in the node list. If the node is in the node list, block 1704 transfers to block 1712. If the node is not in the node list, block 1704 transfers to block 1706 which calls Fig. 13 to determine whether the node is a valid node. If the node is not valid, block 1708 returns to the caller. If the node is valid, block 1708 transfers to block 1710 which calls Fig. 15 to add the node to the node list. After adding the node to the node list, or if the node already existed, control transfers block 1712 which updates the state information about the node in the node list before returning to the caller.

In addition to querying nodes on the network, the discovery system can also query any network probes that may be attached to the network. Information about other nodes on the network can be obtained from these probes, and the discovery system can use this information to assist in discovering other nodes on the network.

Having thus described a presently preferred embodiment of the present invention, it will now be appreciated that the objects of the invention have been fully achieved, and it will be understood by those skilled in the art that many changes in construction and circuitry and widely differing embodiments and applications of the invention will suggest themselves without departing from the spirit and scope of the present invention. The disclosures and the description herein are intended to be illustrative and are not in any sense limiting of the invention, more preferably defined in scope by the following claims.

Claims

1. A computer network node discovery process (120) for determining nodes (206, 216, 218) connected to a computer network (118), said process (120) comprising the steps of:
   (a) obtaining (306), from one node of a set of known nodes on said computer network (118), a list of addresses of one or more other nodes
   (b) repeating step (a) for each of said other nodes obtained; and
   (c) storing said list of node addresses in a file (808); whereby said list of node addresses may be displayed to a user of said computer network.

2. The process of claim 1 further comprising the step of:
   (d) repeating steps (a) through (c) at regular time intervals.

3. The process of claim 2 further comprising the step of:
   (a1) obtaining from each bridge unit (208) connected to said network (118) a list of addresses of all nodes accessible by said bridge unit (208).

4. The process of claim 3 further comprising the step of:
   (a2) obtaining from each router unit (220) connected to said network (118) a list of addresses of all nodes accessible by said router unit (220).

5. The process of claim 4 further comprising the step of:
   (a3) obtaining from each gateway unit (220) connected to said network (118) a list of addresses of all nodes accessible by said gateway unit (220).

6. The process of claim 5 further comprising the step of:
   (a4) obtaining from any network probe device (224) connected to said network (118) a list of addresses of all nodes known to said network probe device (224).

7. A computer network node discovery process (120) for determining nodes connected to a TCP/IP computer network (118), said process comprising the steps of:
   (a) obtaining (306), from one node of a set of known nodes on said computer network, an address translation table containing a list of addresses of other nodes with which said one node communicates;
   (b) repeating step (a) for each of said other nodes in said address translation table;
   (c) storing said list of nodes in a file (808); and
   (d) repeating steps (a) through (c) at regular time intervals.

8. The process of claim 7 further comprising the steps of:

nected to said network (118) an address translation table containing a list of addresses of nodes accessible from said bridge unit (208);

(a2) obtaining from each router unit (220) connected to said network (118) an address translation table containing a list of addresses of nodes accessible from said router unit (220);

(a3) obtaining from each gateway unit (220) connected to said network (118) an address translation table containing a list of addresses of nodes accessible from said gateway unit (220);

(a4) obtaining from any network probe devices (224) attached to said network (118) a list of addresses of all nodes known to said network probe (224); and

(a5) obtaining from each node in said network (118) an interface table and an internet protocol table which defines other networks and nodes to which said node is connected.

9. A computer network node discovery process (120) for determining nodes connected to a computer network (118), said process comprising the steps of:

(a) sending a general response message (307) to all nodes on said network;

(b) creating a node list (410) containing the address of each node responding to said general response message;

(c) obtaining (306), from each node in said node list, a second list of addresses of other nodes with which said node communicates;

(d) adding each node (410) in said second list to said node list;

(e) repeating steps (c) through (d) for each of said nodes in said second list;

(f) storing said node list in a file (808); and

(g) repeating steps (a) through (f) at regular time intervals.

10. The process of claim 9 further comprising the steps of:

(c1) obtaining from each bridge unit (208) connected to said network (118) a list of addresses of all nodes accessible by said bridge unit (208);

(c2) obtaining from each router unit (220) connected to said network (118) a list of addresses of all nodes accessible by said router unit (220);

(c3) obtaining from each gateway unit (220) connected to said network (118) a list of addresses of all nodes accessible by said gateway unit (220); and

(c4) obtaining from any network probe devices (224) attached to said network (118) a list of addresses of all nodes known to the network probe (224).

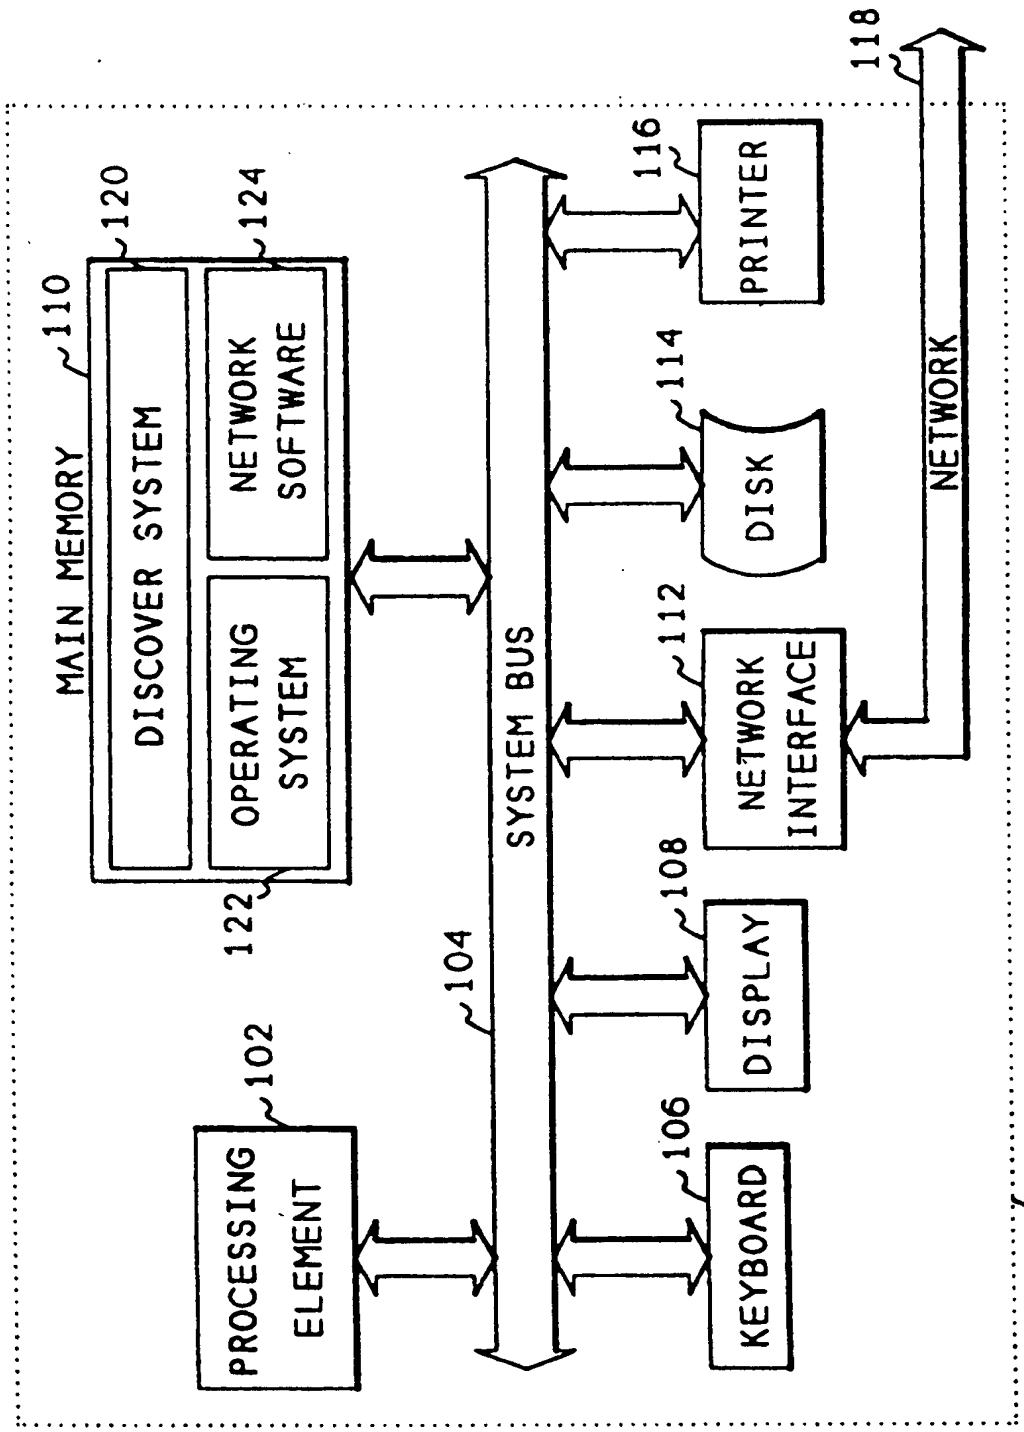
100
PROCESSING ELEMENT
102
104
SYSTEM BUS
MAIN MEMORY
110
DISCOVER SYSTEM
120
NETWORK SOFTWARE
124
OPERATING SYSTEM
122
KEYBOARD
106
DISPLAY
108
NETWORK INTERFACE
112
DISK
114
PRINTER
116
NETWORK
118

FIG. 1

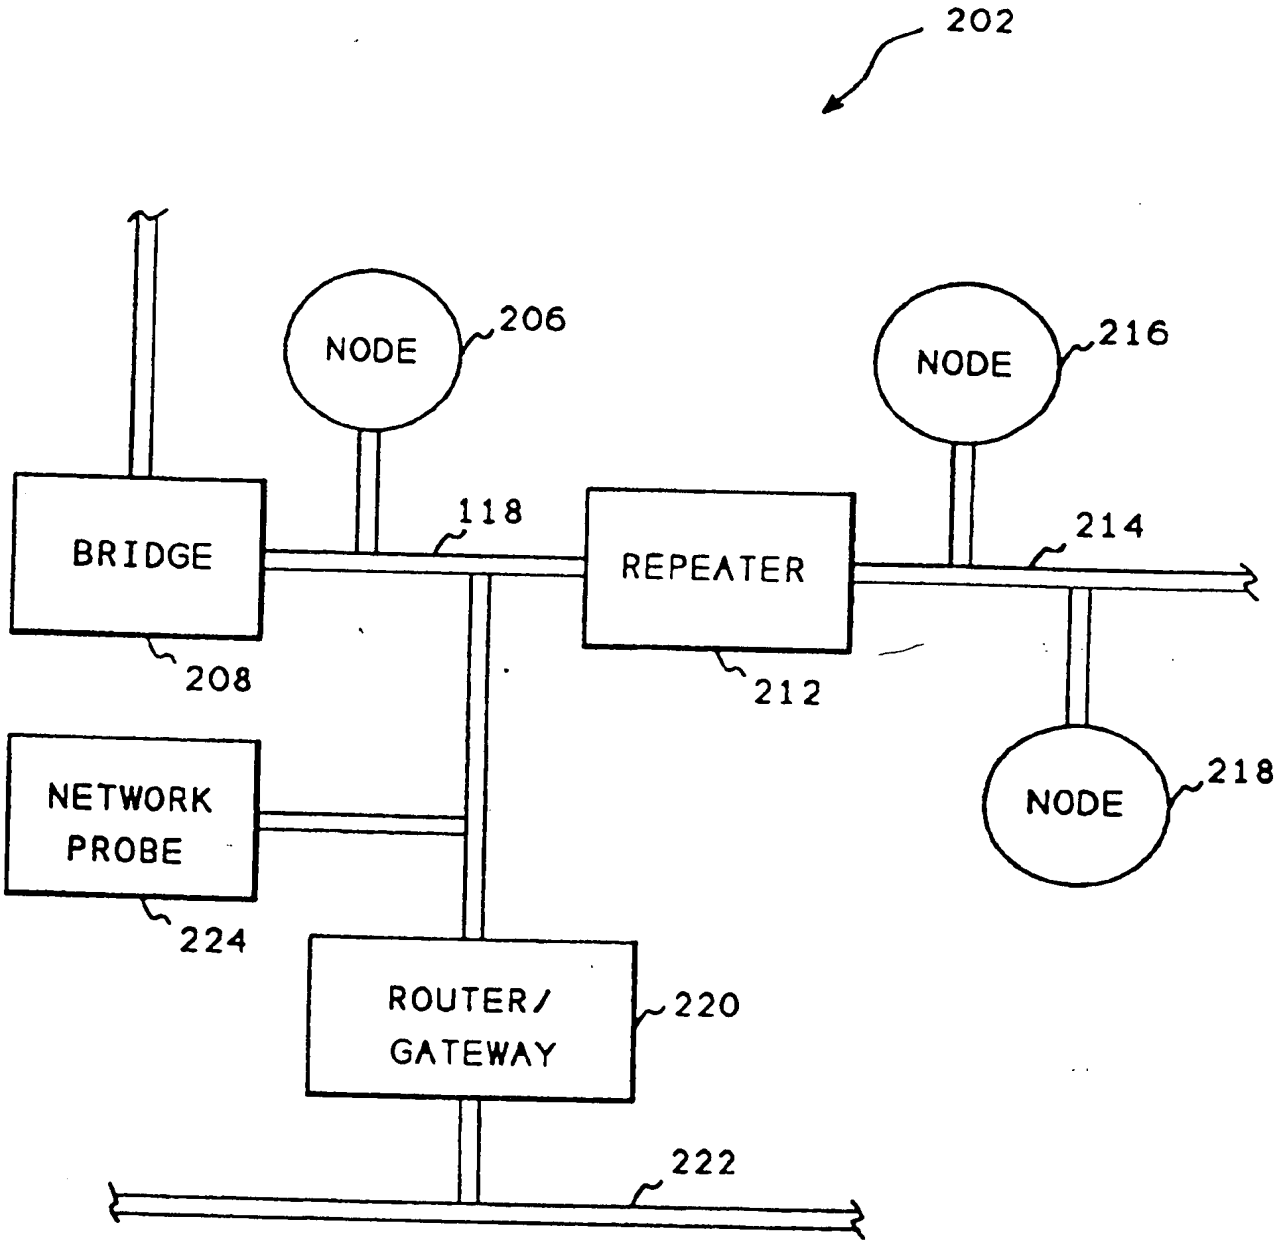
202
NODE
206
NODE
216
BRIDGE
118
REPEATER
214
208
212
NETWORK
PROBE
NODE
218
224
ROUTER/
GATEWAY
220
222

FIG. 2

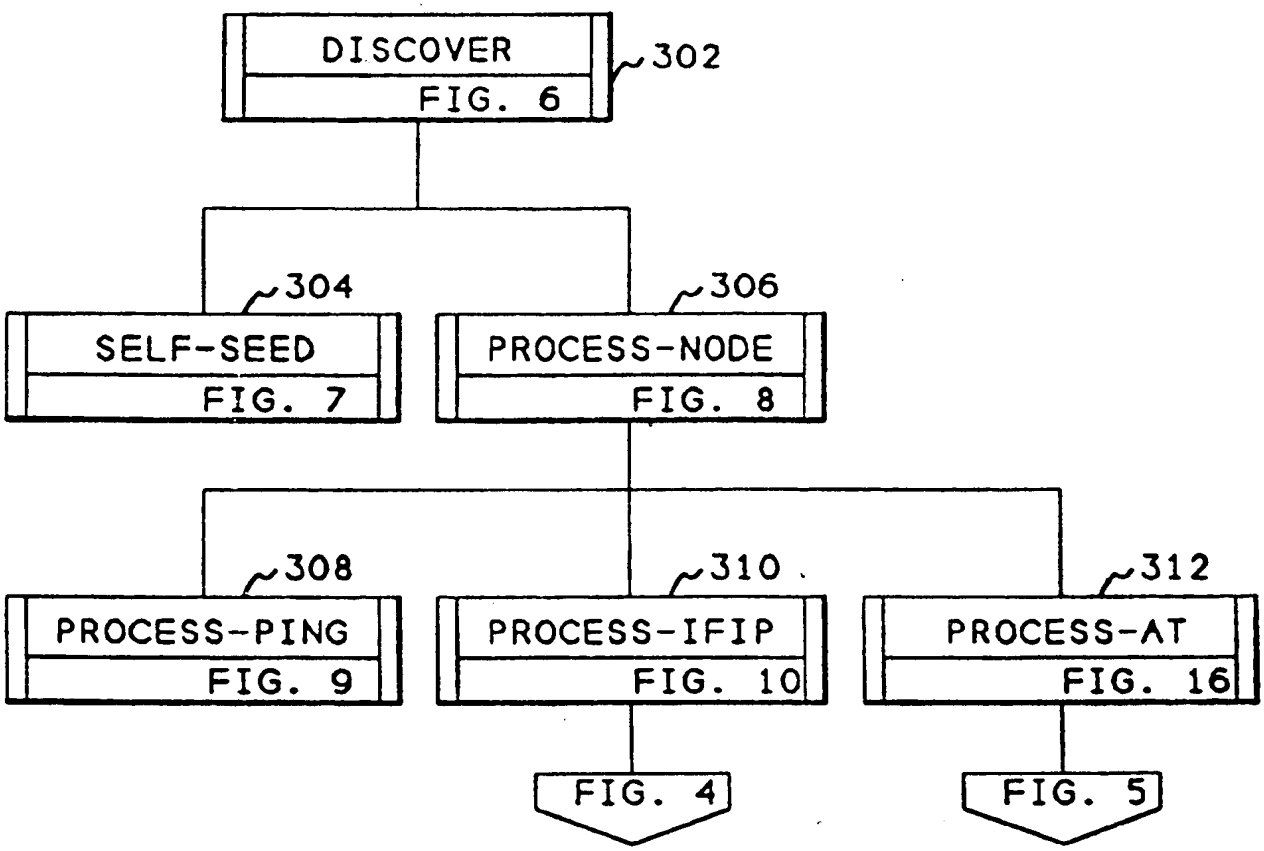
DISCOVER
FIG. 6
302
304
SELF-SEED
FIG. 7
306
PROCESS-NODE
FIG. 8
308
PROCESS-PING
FIG. 9
310
PROCESS-IFIP
FIG. 10
312
PROCESS-AT
FIG. 16
FIG. 4
FIG. 5

FIG. 3

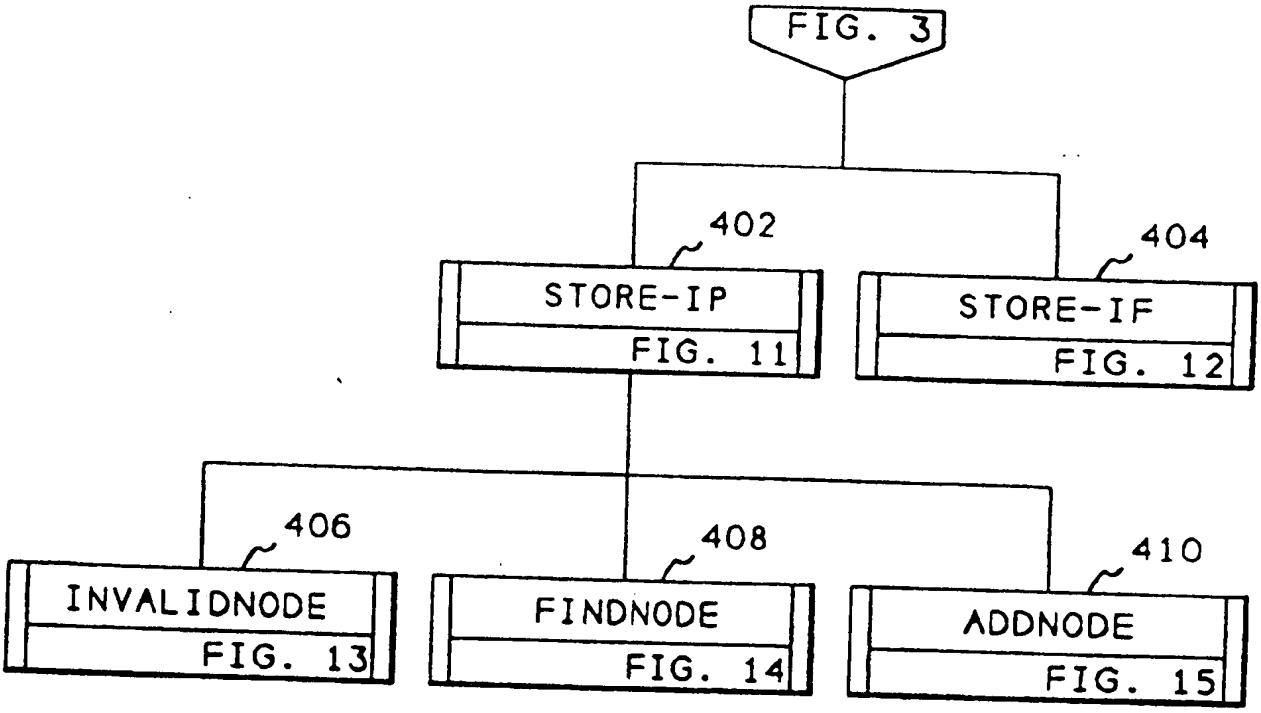
FIG. 3
402
STORE-IP
FIG. 11
404
STORE-IF
FIG. 12
406
INVALIDNODE
FIG. 13
408
FINDNODE
FIG. 14
410
ADDNODE
FIG. 15

FIG. 4

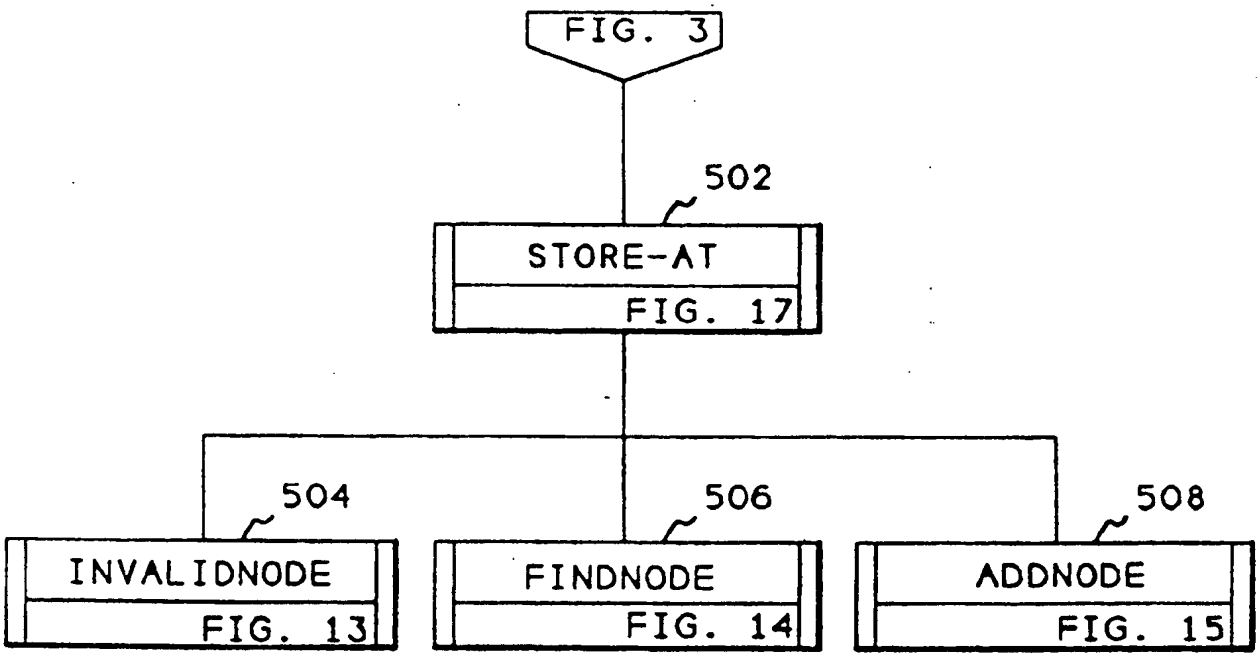
FIG. 3
502
STORE-AT
FIG. 17
504
INVALIDNODE
FIG. 13
506
FINDNODE
FIG. 14
508
ADDNODE
FIG. 15

FIG. 5

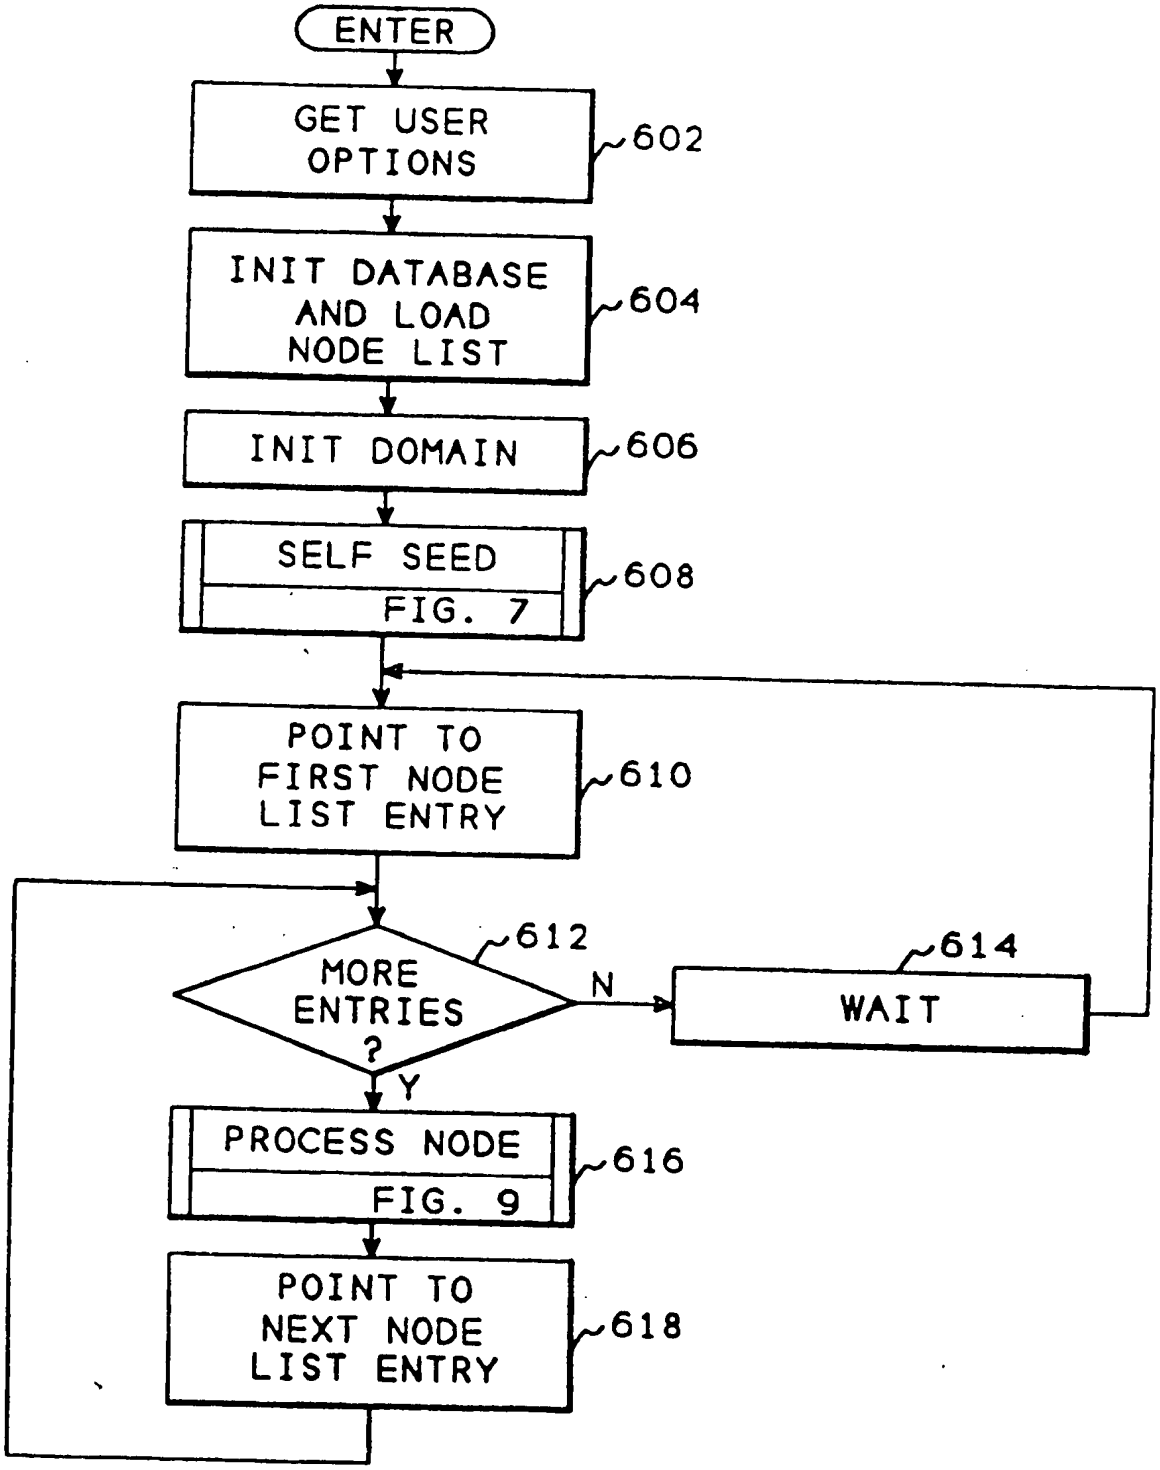
ENTER
GET USER OPTIONS
602
INIT DATABASE AND LOAD NODE LIST
604
INIT DOMAIN
606
SELF SEED
FIG. 7
608
POINT TO FIRST NODE LIST ENTRY
610
612
MORE ENTRIES ?
N
614
WAIT
Y
PROCESS NODE
FIG. 9
616
POINT TO NEXT NODE LIST ENTRY
618

FIG. 6

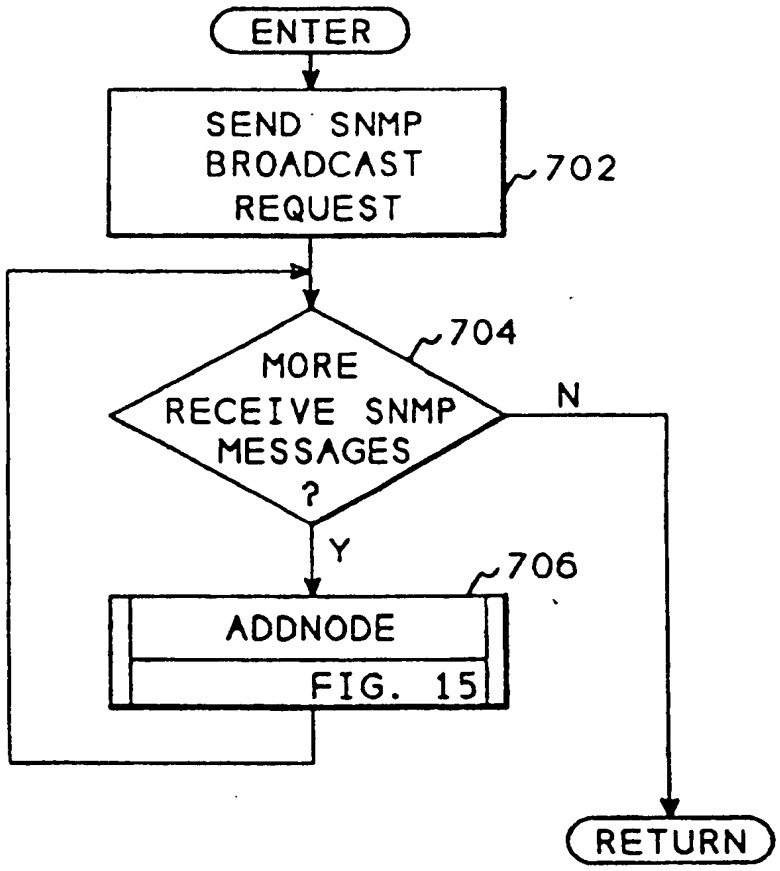
ENTER
SEND SNMP BROADCAST REQUEST
702
MORE RECEIVE SNMP MESSAGES ?
704
N
Y
ADDNODE
FIG. 15
706
RETURN

FIG. 7

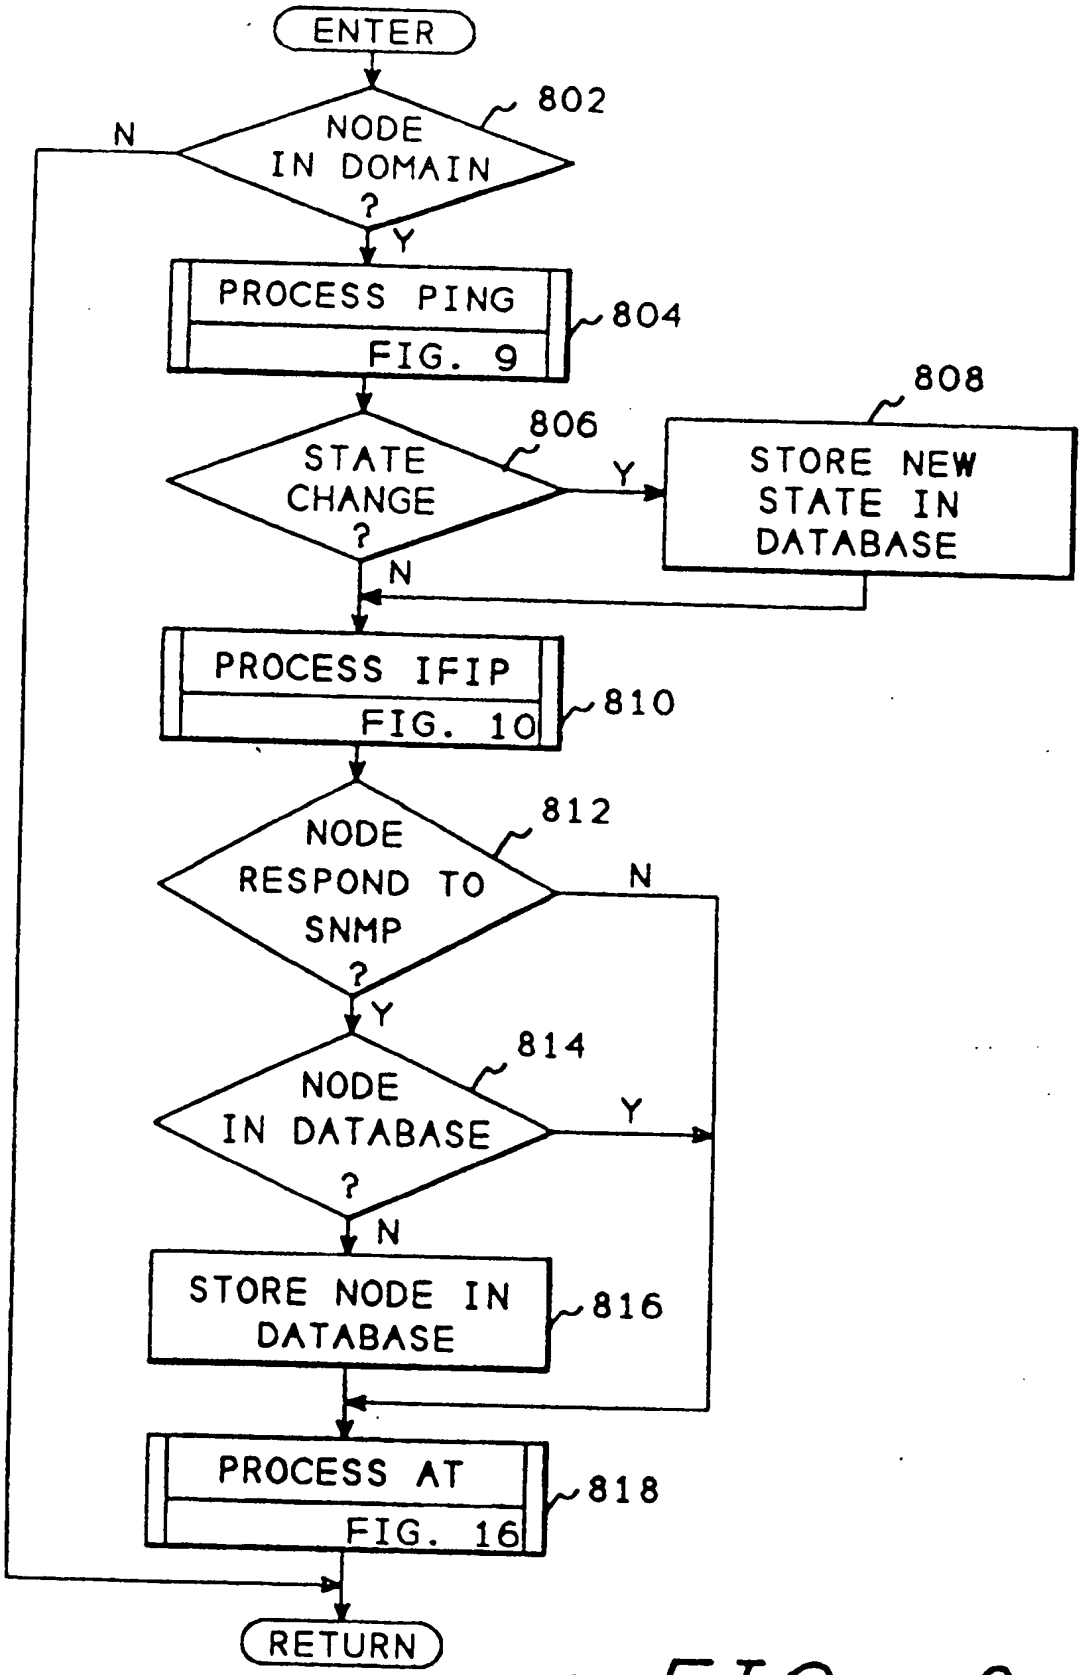
ENTER
802
NODE IN DOMAIN ?
N
Y
PROCESS PING
FIG. 9
804
808
806
STATE CHANGE ?
Y
STORE NEW STATE IN DATABASE
N
PROCESS IFIP
FIG. 10
810
812
NODE RESPOND TO SNMP ?
N
Y
814
NODE IN DATABASE ?
Y
N
STORE NODE IN DATABASE
816
PROCESS AT
FIG. 16
818
RETURN

FIG. 8

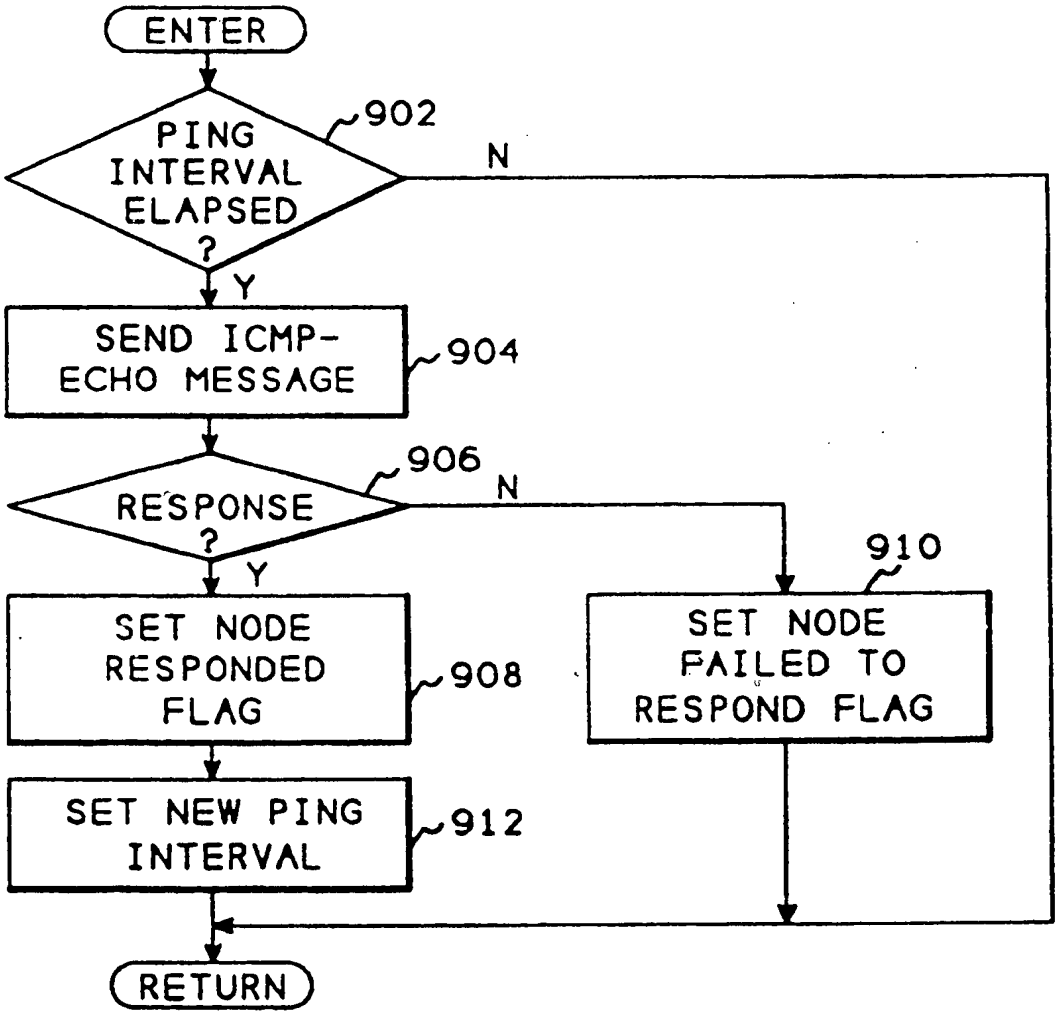
ENTER
PING INTERVAL ELAPSED ?
902
N
Y
SEND ICMP-ECHO MESSAGE
904
RESPONSE ?
906
N
Y
SET NODE RESPONDED FLAG
908
SET NODE FAILED TO RESPOND FLAG
910
SET NEW PING INTERVAL
912
RETURN

FIG. 9

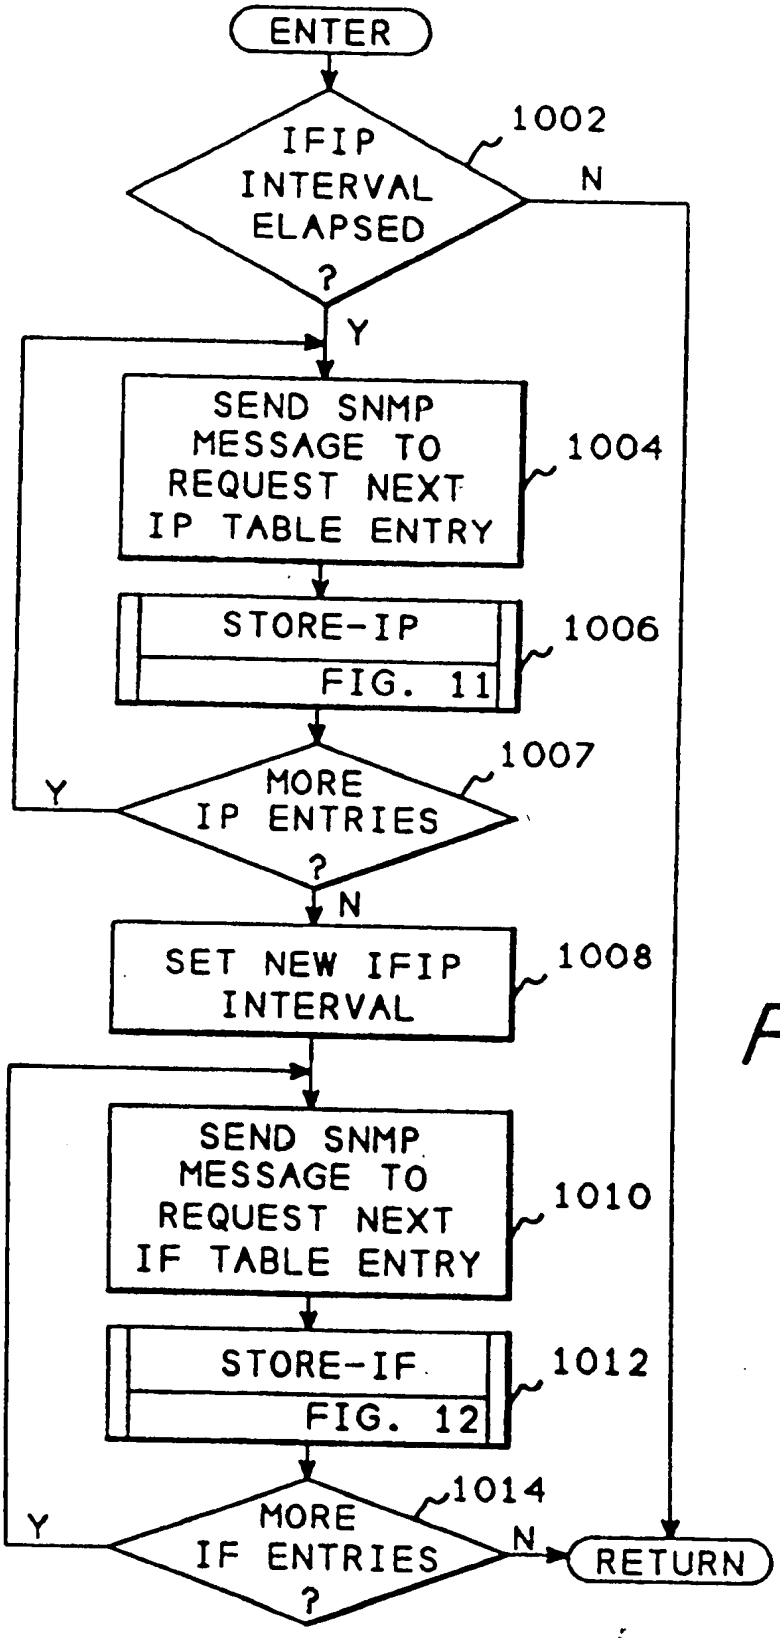
ENTER
IFIP INTERVAL ELAPSED ?
1002
N
Y
SEND SNMP MESSAGE TO REQUEST NEXT IP TABLE ENTRY
1004
STORE-IP
FIG. 11
1006
MORE IP ENTRIES ?
1007
Y
N
SET NEW IFIP INTERVAL
1008
SEND SNMP MESSAGE TO REQUEST NEXT IF TABLE ENTRY
1010
STORE-IF
FIG. 12
1012
MORE IF ENTRIES ?
1014
Y
N
RETURN

FIG. 10

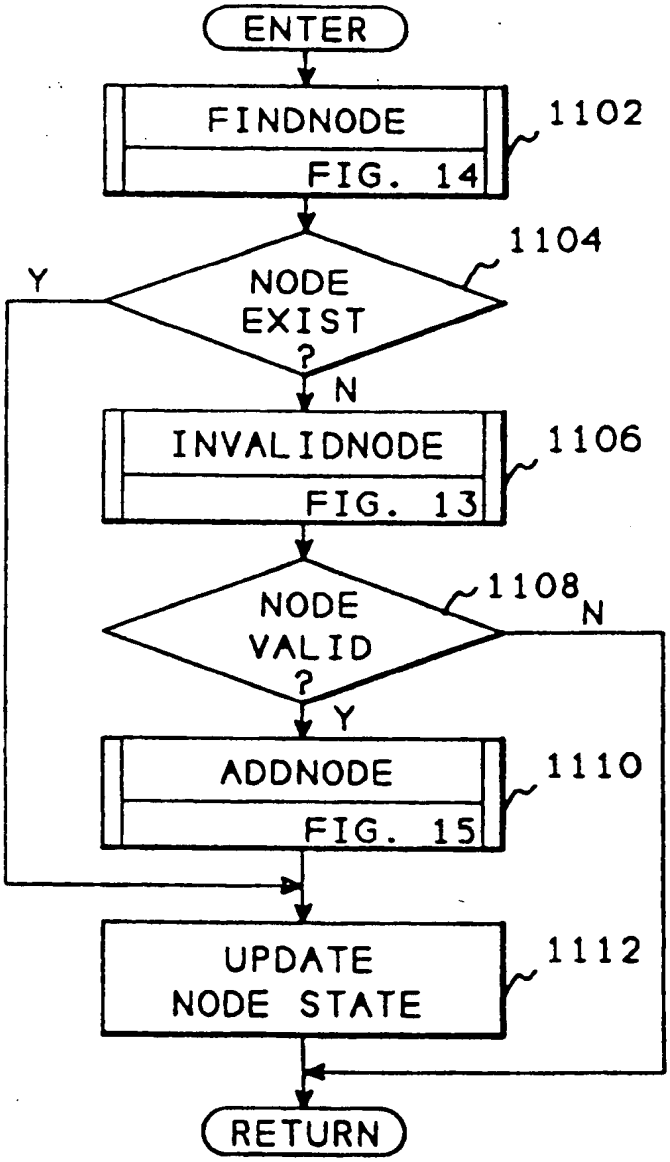
ENTER
FINDNODE
FIG. 14
1102
1104
NODE EXIST ?
Y
N
INVALIDNODE
FIG. 13
1106
NODE VALID ?
1108
N
Y
ADDNODE
FIG. 15
1110
UPDATE NODE STATE
1112
RETURN

FIG. 11

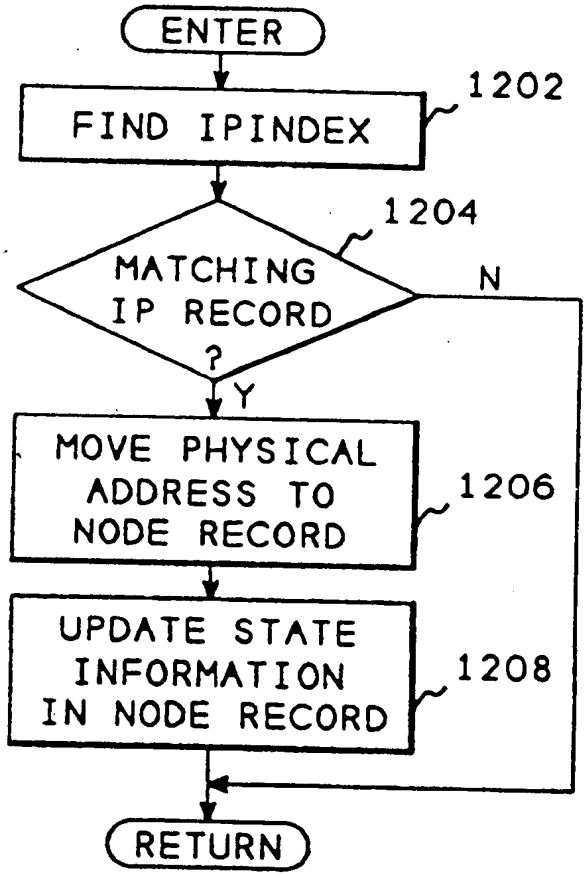
ENTER
FIND IPINDEX
1202
MATCHING IP RECORD ?
1204
N
Y
MOVE PHYSICAL ADDRESS TO NODE RECORD
1206
UPDATE STATE INFORMATION IN NODE RECORD
1208
RETURN

FIG. 12

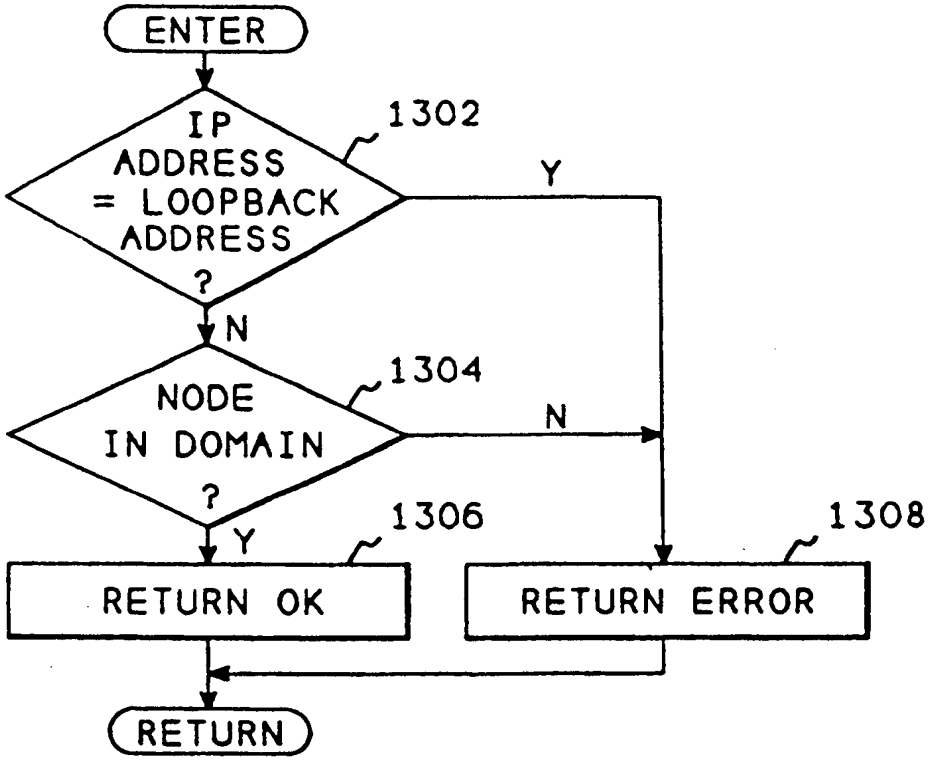
ENTER
IP ADDRESS = LOOPBACK ADDRESS ?
1302
Y
N
NODE IN DOMAIN ?
1304
N
Y
1306
1308
RETURN OK
RETURN ERROR
RETURN

FIG. 13

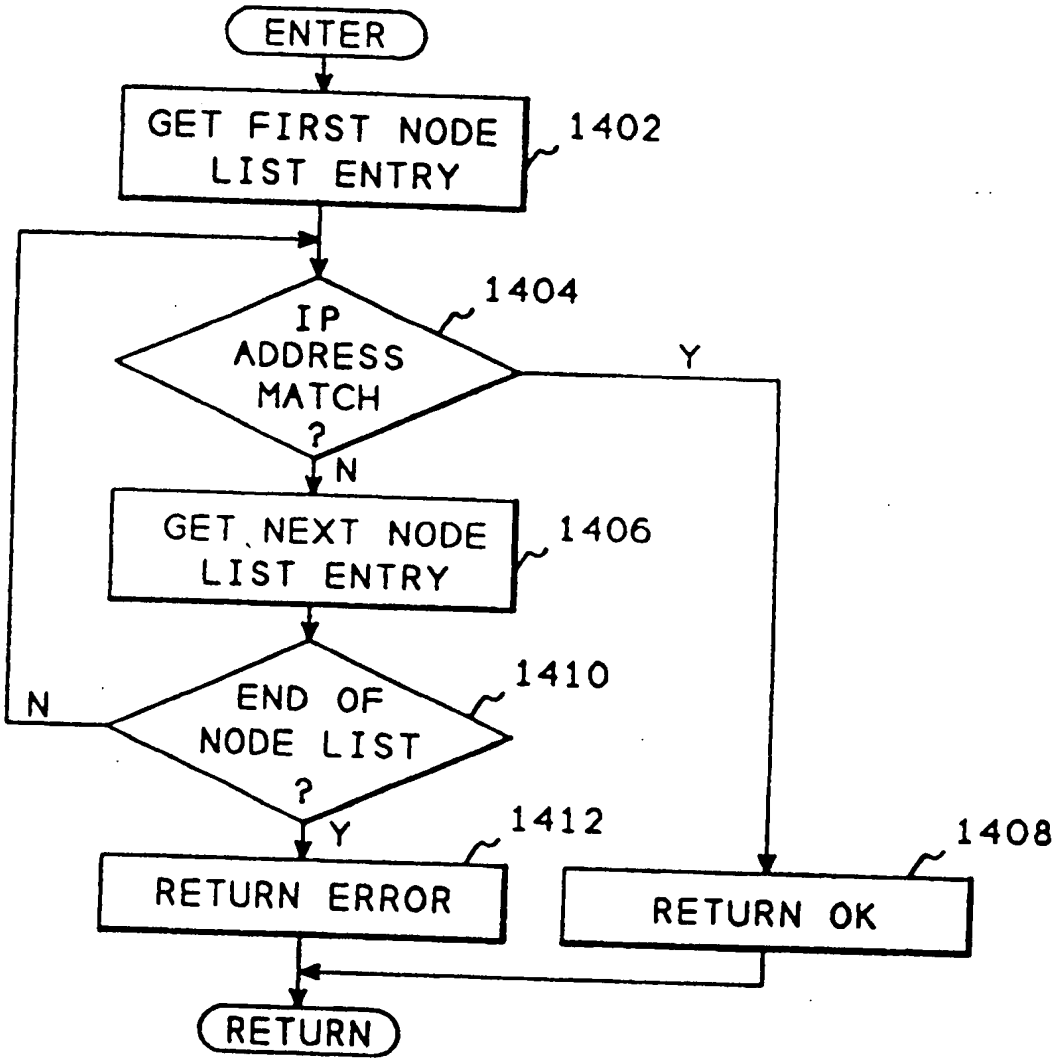
ENTER
GET FIRST NODE LIST ENTRY
1402
IP ADDRESS MATCH ?
1404
Y
N
GET NEXT NODE LIST ENTRY
1406
END OF NODE LIST ?
1410
N
Y
RETURN ERROR
1412
RETURN OK
1408
RETURN

FIG. 14

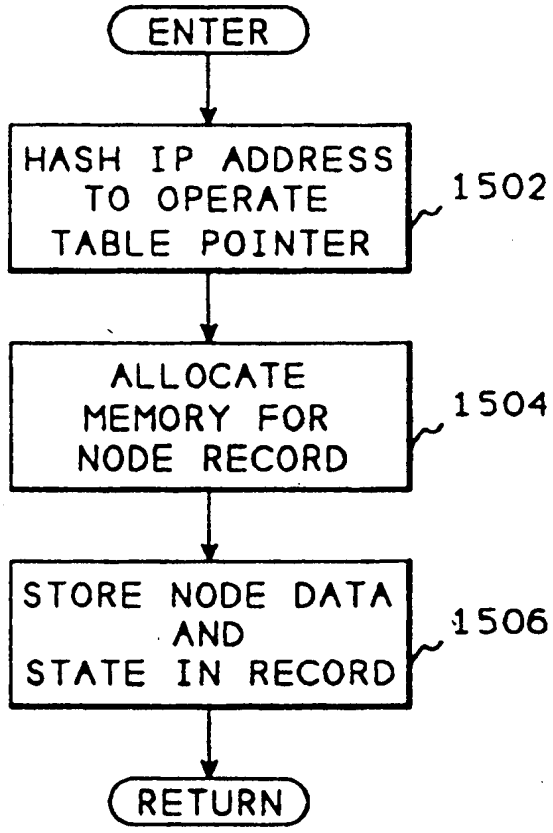
ENTER
HASH IP ADDRESS TO OPERATE TABLE POINTER
1502
ALLOCATE MEMORY FOR NODE RECORD
1504
STORE NODE DATA AND STATE IN RECORD
1506
RETURN

FIG. 15

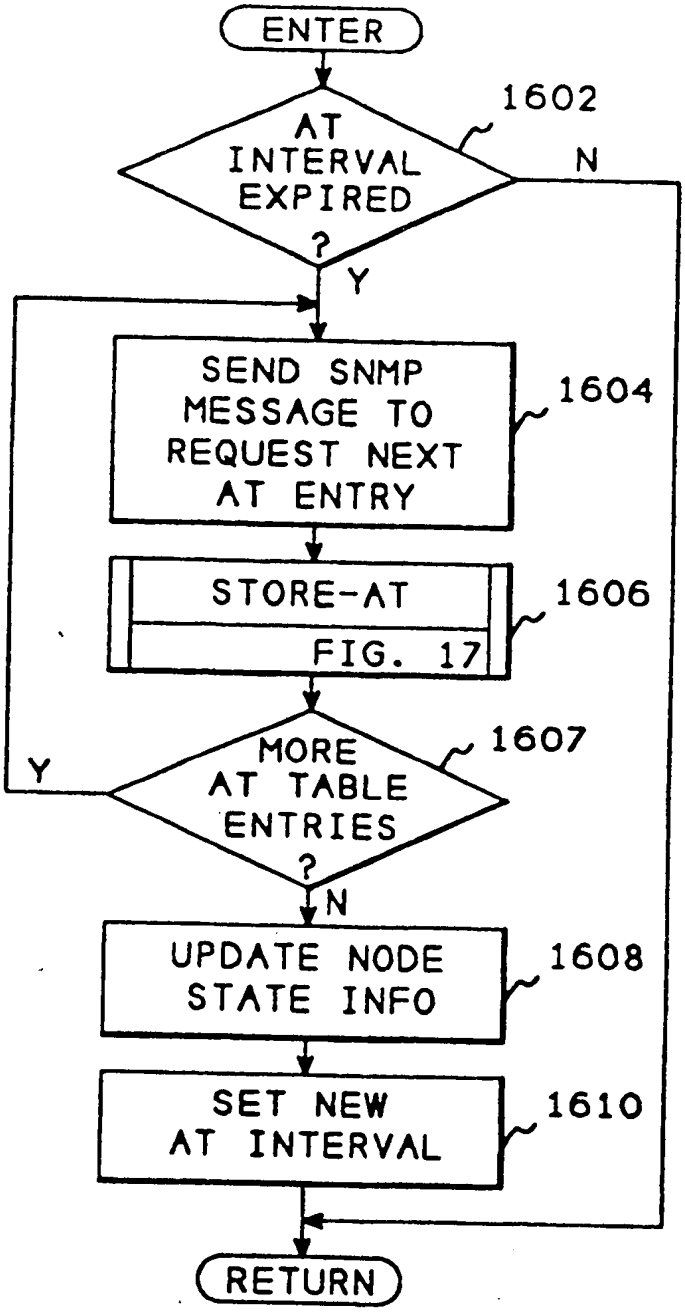
ENTER
AT INTERVAL EXPIRED ?
1602
N
Y
SEND SNMP MESSAGE TO REQUEST NEXT AT ENTRY
1604
STORE-AT
FIG. 17
1606
MORE AT TABLE ENTRIES ?
1607
Y
N
UPDATE NODE STATE INFO
1608
SET NEW AT INTERVAL
1610
RETURN

FIG. 16

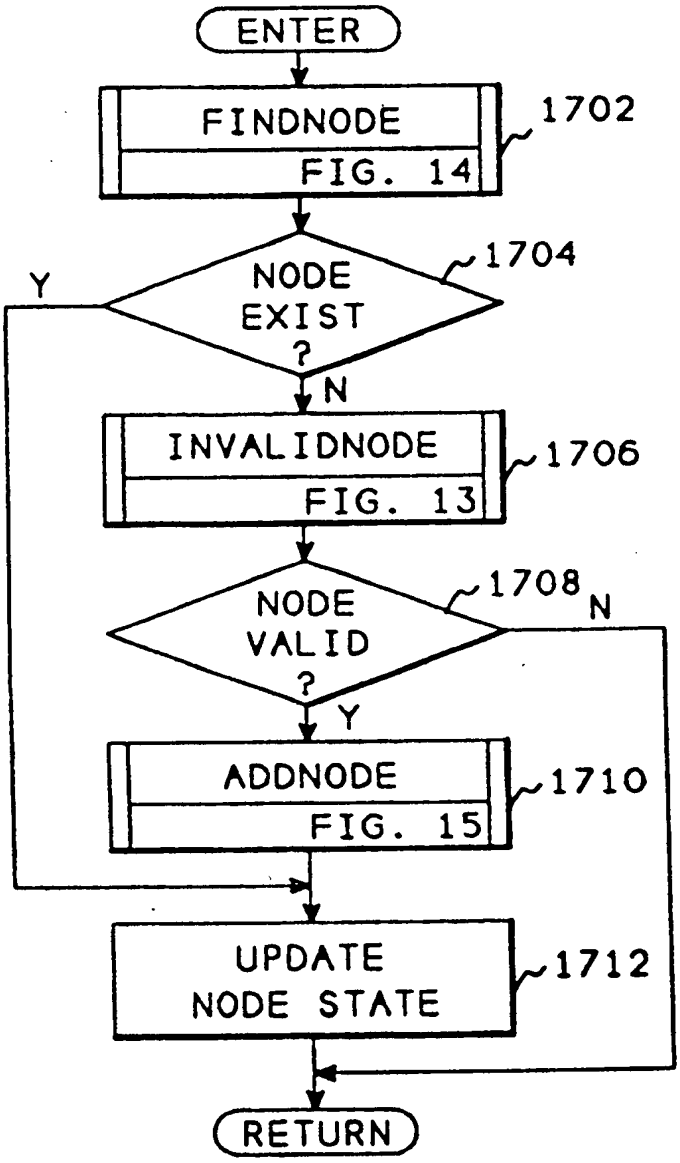
ENTER
FINDNODE
FIG. 14
1702
NODE
EXIST
?
1704
Y
N
INVALIDNODE
FIG. 13
1706
NODE
VALID
?
1708
N
Y
ADDNODE
FIG. 15
1710
UPDATE
NODE STATE
1712
RETURN

FIG. 17